

# UTILITY PATENT APPLICATION TRANSMITTAL

## (Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.  
20496/2

Total Pages in this Submission

### TO THE ASSISTANT COMMISSIONER FOR PATENTS

Box Patent Application  
Washington, D.C. 20231

Transmitted herewith for filing under 35 U.S.C. 111(a) and 37 C.F.R. 1.53(b) is a new utility patent application for an invention entitled:

**METHOD AND APPARATUS FOR PROVISIONING, ASSIGNMENT, MEASUREMENT AND AUDIT SERVICES PROVIDED TO CLUSTERED NETWORK SITE USERS, AND TO GENERATE SERVICE LEVEL AGREEMENT REPORTS**

and invented by:

**Krishna Mangipudi and Vijay Basani**

If a **CONTINUATION APPLICATION**, check appropriate box and supply the requisite information:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: \_\_\_\_\_

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: \_\_\_\_\_

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: \_\_\_\_\_

Enclosed are:

#### Application Elements

1. ☒ Filing fee as calculated and transmitted as described below
2. ☒ Specification having 15 pages and including the following:
  - a. ☒ Descriptive Title of the Invention
  - b. ☐ Cross References to Related Applications (if applicable)
  - c. ☐ Statement Regarding Federally-sponsored Research/Development (if applicable)
  - d. ☐ Reference to Microfiche Appendix (if applicable)
  - e. ☒ Background of the Invention
  - f. ☒ Brief Summary of the Invention
  - g. ☒ Brief Description of the Drawings (if drawings filed)
  - h. ☒ Detailed Description
  - i. ☒ Claim(s) as Classified Below
  - j. ☒ Abstract of the Disclosure

# UTILITY PATENT APPLICATION TRANSMITTAL (Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.  
20496/2

Total Pages in this Submission

## Application Elements (Continued)

3. ☒ Drawing(s) (when necessary as prescribed by 35 USC 113)
- a. ☐ Formal Number of Sheets \_\_\_\_\_
- b. ☒ Informal Number of Sheets 6
4. ☒ Oath or Declaration
- a. ☒ Newly executed (original or copy) ☐ Unexecuted
- b. ☐ Copy from a prior application (37 CFR 1.63(d)) (for continuation/divisional application only)
- c. ☒ With Power of Attorney ☐ Without Power of Attorney
- d. ☐ DELETION OF INVENTOR(S)  
Signed statement attached deleting inventor(s) named in the prior application,  
see 37 C.F.R. 1.63(d)(2) and 1.33(b).
5. ☐ Incorporation By Reference (usable if Box 4b is checked)  
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. ☐ Computer Program in Microfiche (Appendix)
7. ☐ Nucleotide and/or Amino Acid Sequence Submission (if applicable, all must be included)
- a. ☐ Paper Copy
- b. ☐ Computer Readable Copy (identical to computer copy)
- c. ☐ Statement Verifying Identical Paper and Computer Readable Copy

## Accompanying Application Parts

8. ☐ Assignment Papers (cover sheet & document(s))
9. ☐ 37 CFR 3.73(B) Statement (when there is an assignee)
10. ☐ English Translation Document (if applicable)
11. ☐ Information Disclosure Statement/PTO-1449 ☐ Copies of IDS Citations
12. ☐ Preliminary Amendment
13. ☒ Acknowledgment postcard
14. ☒ Certificate of Mailing

☐ First Class ☒ Express Mail (Specify Label No.): EE 633 278 765 US

# UTILITY PATENT APPLICATION TRANSMITTAL (Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.  
20496/2

Total Pages in this Submission

## Accompanying Application Parts (Continued)

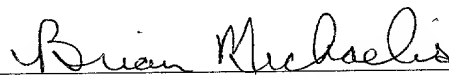
15. ☐ Certified Copy of Priority Document(s) (if foreign priority is claimed)
16. ☐ Additional Enclosures (please identify below):

## Fee Calculation and Transmittal

### CLAIMS AS FILED

For	#Filed	#Allowed	#Extra	Rate	Fee
Total Claims	3	- 20 =	0	x \$18.00	\$0.00
Indep. Claims	1	- 3 =	0	x \$78.00	\$0.00
Multiple Dependent Claims (check if applicable) <input type="checkbox"/>					\$0.00
BASIC FEE					\$760.00
OTHER FEE (specify purpose)					\$0.00
TOTAL FILING FEE					\$760.00

- ☒ A check in the amount of \$760.00 to cover the filing fee is enclosed.
- ☒ The Commissioner is hereby authorized to charge and credit Deposit Account No. 50-0369 as described below. A duplicate copy of this sheet is enclosed.
- ☐ Charge the amount of as filing fee.
- ☒ Credit any overpayment.
- ☒ Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.
- ☐ Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).

  
Signature

Brian L. Michaelis (Reg. No. 34,221)  
Brown Rudnick Freed & Gesmer, PC  
One Financial Street  
Boston, MA 02111

Dated: November 30, 1999

cc:

5 METHOD AND APPARATUS FOR PROVISIONING, ASSIGNMENT,  
MEASUREMENT AND AUDIT SERVICES PROVIDED TO CLUSTERED  
NETWORK SITE USERS, AND TO GENERATE SERVICE LEVEL AGREEMENT  
REPORTS.

#### Related Application

10 Priority is claimed to U.S. Provisional Patent Application Serial No. 60/110350,  
entitled A SOFTWARE SYSTEM THAT PROVIDES A MECHANISM TO  
PROVISION, ASSIGN, MEASURE, & AUDIT SERVICES OFFERED TO  
INTERNET/INTRANET SITE USERS, AND GENERATE QUANTIFIABLE  
REPORTS – SERVICE LEVEL AGREEMENTS (SLA) AND SLA REPORTING, filed  
November 19, 1998, which is incorporated herein by reference.

15

#### Field of the Invention

The present invention pertains to network communications, particularly for  
establishing, measuring, and reporting service attributes.

20

#### Background of the Invention

25 The rate of flow of data in computer networks between hosts and clients in  
Internets and Intranets depends upon many parameters. Some of these parameters can be  
tied to the provision of resources. These provisioned resources can be measured and audit  
reports can be generated to see if the parameters are in the range of negotiated Service  
Level Agreements. A Service Level Agreement (SLA) between a service-provider and a  
user defines the expected and acceptable properties of the services, typically in the  
context of providing Internet services. The SLA provides a tool by which performance  
goals can be measured, by defining the performance metrics and the corresponding goals.  
30 By monitoring compliance with SLA limits, a service provider can avoid the costly  
problems that result from disappointing users or hosted customers.

09451592 13069

Network operations can be monitored and measured using standard techniques such as Route Monitor (RMON) and its probes to gain insight into the flow rates of data between points within these monitored networks. These measurements stop short of the application layer in the OSI model. Application layer parameters such as throughput, latency and round trip time are not covered in these measurements. Other factors that influence the round trip time at the application layer are local conditions such as CPU availability (processing overload), and secondary resource availability (e.g., database access). Furthermore, the known network monitors do not monitor the number of concurrent network connections that can be opened on each server. A web site on the Internet or Intranet may contain numerous, diverse servers, each with its own CPU, databases, and network connections. Thus, network layer measurements only shed partial light on the performance of a web site.

It is known that a SLA can be defined to guarantee the flow rates in networks and these SLAs can be honored in switched networks using such protocols as Reservation Protocol (RSVP) or in the ATM fabric at a rather coarse granular level. Network bandwidth is then assigned to the flows based on the SLA parameters. This SLA-based assignment guarantees the requested bandwidth from the client to the web server and back. However, it stops short of measuring the traffic flow up to the application layer at the web server that provides the service. In the context of the application layer (OSI layer 7) in the HyperText Transfer Protocol (HTTP) – as it pertains to the flows in the Internet – there are several parameters that can be provisioned (i.e., installed and activated) and then measured and audited. In order to guarantee end-to-end SLA these parameters have to be taken into account by the monitoring system.

It is known by those skilled in the art that individual host computers can create logs of each client request. These log files are stored, usually in ASCII format on disk in the host computers. The log files contain “raw,” unformatted information about each transaction or client request, and may be provided in diverse, incompatible formats. Further, as mentioned above, these log files contain only a part of the information necessary to generate reports about SLAs.

Within a cluster of web servers there is often an autonomous sharing of resources to service an external request more efficiently. Simple network performance monitoring

reports or host performance monitoring reports do not collect and correlate information in ways that can assist in evaluating and targeting network elements that may cause violations of an SLA. Even if host performance and network performance reports are combined, existing tools do not provide a way to filter out reports of problems that are automatically handled by other systems (e.g., automatic retry). One major disadvantage of the prior art is the inability to monitor and characterize real-time request streams and their corresponding responses. Another disadvantage is the inability to match the measured parameters with each independent SLA in a manner that provides user-oriented reporting. Yet another disadvantage is that existing reporting mechanisms are necessarily tied to particular machines, even though a user transaction may be serviced by any of several different machines. Similarly, reporting on the performance related to some particular web content (e.g., a web site) is difficult when the same content can be served by any one of several different machines.

One example of a known SLA implementation is disclosed in U.S. Pat. No. 5,893,905, issued April 13, 1999. In that system, as applied to a scheduled computer processing job environment, a monitoring center automatically retrieves job exception data (logs), job run data, and clocktime data from multiple computer systems, each of which is running a collection and retrieval program module. The retrieved data is stored in appropriate databases for each type of data collected. A jobflow table, according to the daily SLAs, is also stored in the system, corresponding to a set of tasks to be performed. A "periodic data analysis process" determines whether jobs are timely run, or if errors have occurred. If tardy jobs or errors are detected, the system determines whether the result will negatively affect the SLA. If a problem is detected, then operators are signaled with reports designating jobs that may impact an SLA, and which SLA is in jeopardy, so that operations personnel can take additional manual steps.

One major disadvantage of the disclosed system is the reliance upon pre-defined SLA jobflow tables for determining which jobs should be run at a given time on a given day. The jobflow tables presume a static jobflow. The tables also presume a predictable timing, either for a job, or for a given series of jobs necessary to comply with an SLA. Furthermore, the disclosed system provides an alert only if a job error has occurred or if estimated time to complete a job exceeds the limits of the corresponding SLA. The only

information obtained is that the schedule of a downstream job may be affected. These limited signals cannot be easily correlated with the wide variety of metrics that can have real-time affect upon users. A static job table cannot be applied in the environment of a real-time web-sever where there is no standardized sequence of jobs, and "time of day" sequencing is irrelevant. Nor can this type of limited output signaling be used to determine whether a problem is temporary or persistent. Also, the limited output of the prior art system does not accommodate reporting on multiple "back-end" servers that can share the role of servicing real-time requests; rather, it simply reports a "violation." The same report would be issued even if the job were re-run on another production server.

### Summary of the Invention

The present invention provides a method and apparatus for generating, collecting, and manipulating useful information for validating or defining SLAs of web servers on a network. Web servers comprising a web farm on the network are adapted for logging detailed runtime information regarding user transactions and performance parameters.

According to the invention, an Accumulator device on the network interacts with an intelligent agent on each web server to collect and combine their log files, process the combined file and post information into a database. An operator enters committed performance parameters into an SLA Reporter system according to classes of users, URLs, transactions, content or file type, or classes of web sites being hosted on the web servers. When compared with the actual data from the database, processing of SLA reports indicate how well the parameters of the SLAs are being met for users, virtual sites, classes, URL's and transactions, or other measurable elements. By generating, collecting, combining and processing in this manner, application-specific performance can be quickly and automatically evaluated with respect to parameters related to user satisfaction and detailed signals can be issued for cases in which remedial steps should be undertaken.

In further accord with the present invention, a network model includes a plurality of computer systems clustered into a web farm, in which a front-end system distributes service requests to one or more back-end web servers. The front-end system receives a service request from a user outside the web farm, selects one of the back-end servers to

service the request, and forwards (routes) the request to the selected back-end server. Each back-end server then services the user request, and generates any necessary response directly to the requesting user. A given back-end server may be requested to process transactions destined for any of a number of "virtual sites" that the back-end  
5 server is hosting as a surrogate for the "hosted site" addressed in the transaction.

Each back-end server generates its own time-stamped log entry for each transaction it processes, resulting in a distributed log file in the network. The individual files of the distributed log file each include information regarding username, date, time, service, server name, server IP address, client IP address, processing time, bytes sent,  
10 bytes received, service status, operation, target URL, User Agent, referrer parameters, Smsserver, Smvirtual site, and cookie. Log files collected in the web farm also include a time stamp to measure the web-farm transfer time, the front end server's name, and the DNS name of the host that serves the request. This extensive combination of information has the advantage of containing the specific details necessary for building a meaningful  
15 SLA report. A further advantage is realized by including user, server, URL, and other labels in the log files such that transactions can be traced from one place in the web farm to another, and the corresponding performance elements aggregated. Similarly, detailed logfile labeling facilitates later correlation of reports along any logged parameter, such as according to the URL.

20 An Accumulator process working in conjunction with the intelligent agents on the servers periodically collects the most recent entries of each distributed log file. According to the invention, a periodic event triggers the collection process whereby each back-end server on a list of servers is contacted and a log file update is requested from each server. An agent on each back-end server that is contacted then determines which  
25 recent log information should be transferred to the Accumulator, opens a network transfer connection to the Accumulator, and sends the requested log file update. The Accumulator waits until it has received log file updates from each listed agent and then combines the collected log files into a consolidated log file for the listed back-end servers. The Accumulator may perform the same collection operation for different sets of  
30 servers and combine the log files into different consolidated log files for each set of



03451589 113099  
servers. The combined log files are then parsed and the information is deposited into an SLA database.

An SLA report generator, according to the invention, is configured with the service parameters corresponding to each SLA of a user, or a class of users, a hosted site, or a class of hosted sites or classes of URLs, transactions, or content or file type. By accessing the SLA database, built from information obtained in the combined logs, the SLA report generator can compare the service parameters of each SLA with the actual service delivered, according to the information captured in the distributed log files of the back-end servers.

Some of the advantages of the invention include the following. There is no pre-defined job schedule necessary for comparing performance. Any transaction can be routed to any back-end server without having to adjust the SLA service parameters. Many different kinds of transactions can be handled rather than having to obtain specific job-completion schedules. Unlike the prior art, a system adapted according to the invention provides performance information about not only the individual server processing requests, but also information about the specific machine selected to perform the service, and each major delay element encountered by a transaction. Furthermore, it provides detailed information about the specific "virtual host" to which the request pertained. It also allows for measurement of SLA metrics based on URLs, transactions, or content or file type. Another advantage of the invention is that different classes of users, classes of transactions or URLs, or hosted sites are defined, allowing reports to be prepared on a per-class basis rather than on simply a customer or job basis. Furthermore, a back-end server can be dynamically reconfigured to serve as a virtual host of another class and yet have its log file entries collected and processed properly for each listed class during the proper time.

#### Brief Description of the Drawings

These and other features of the present invention will be better understood in view of the following detailed description taken in conjunction with the drawings, in which:

Fig .1 is a functional block diagram of a network system that implements an embodiment of the invention;

Fig. 2 illustrates the system architecture for accumulating reports from agents;  
Fig. 3 illustrates an embodiment of a report Accumulator;  
Fig. 4 illustrates an embodiment of an SLA report generator; and  
Figs. 5a-5g are illustrations of SLA reports generated according to the invention.

5

### Detailed Description of the Invention

Networks of computer systems have evolved to include those computers and networks that operate under the *de facto* standards of the Internet Protocols (IP). A so-called "web browser" in a client computer 102 permits computer users to access the internet 104 and all of the compatible web sites, as shown in Fig. 1. A user typically enters the name of a web site 118 (the domain name) into the web browser on the user's (client) computer 102. The web browser then uses the Internet services to query a Domain Name Server 106 to locate the address of a web server which hosts the web site, and to receive an Internet address for the web server. The user's web browser then uses the Internet services to contact the web server 108 at the designated Internet address with transactions for the web site 118.

The web site 118 may be hosted on a web server that is available on a single computer that will respond to the transaction itself (e.g., a request, or information), or it may be a distributed system of multiple computers, such as a "web farm" comprised of a front-end sever 108 and multiple back-end servers 210. Multiple back-end servers are typically connected to a local area network (LAN) 110 (or across a wide-area network) which is also connected to the front-end server 108. The use of multiple computers is indicated for reasons such as load balancing, increased availability, security, or other policy/service considerations. However, a typical user is not concerned with the details of where or how the web server 108 handles the web browser interactions with the web site 118, except when performance or availability affect the user directly.

The subscribers of a web server service include information creators (web-site content providers) and others who desire to offer web services to users (e.g. "e-tailers," or merchants, customer service organizations, corporate organizations, as well as the web-site-hosting customers of an ISP, etc). The subscribers can be enticed to enter contracts for services with web server operators by negotiation of a Service-Level Agreement

66081125460

(SLA). An SLA defines the operational parameters of the web server that the operators will monitor, and the relative acceptability of the server services for each parameter. For example, a subscriber may demand a performance objective in which a specified bandwidth (e.g., octets of data per minute), or response time or error rate must be sustained by the server. The operators also need to capture data regarding the actual usage of the server (e.g., for billing purposes), while also predicting the need for additional resources. Failure to adequately provision a web server with the necessary resources (disk drives, processors, communication links, etc) can leave an operator in violation of some parameters of an SLA. On the other hand, a web server that exceeds the provisions of the SLA can also result in rewards for the server-site operator.

In Fig. 1 the system that goes by the DNS (web site) name `www.webmanage.com` is the front-end system 108. "Hosted Sites" are the published domain names which have an IP address 116 of the front-end server 108. By definition, hosted sites do not serve content; however, they forward transactions to the back-end servers 210 in the web farm. "Virtual Sites" are the sites hosted on the back-end web Servers. In Fig. 1, `Ws1.webmanage.com` and `Ws2.webmanage.com` are virtual sites. Virtual Sites are tied to web servers and are responsible for generating responses to user transaction requests. For every hosted site there has to be at least one Virtual Site. Multiple hosted sites may be hosted on the same server and thus have the same IP or Internet address 116. Requests for any of the multiple hosted sites may therefore arrive at the same front end server 108 for multiple hosted sites.

In an embodiment of the present invention, the front end server 108 receives a user's transaction information, such as a request to download particular information from the hosted site 118 located at a particular Internet address 116. The web farm management can choose which back-end web server 210A, 210B, 210C, 210D to use for each hosted web site, and can change servers as a function of load-balancing techniques among the front end server 108 and the back end servers 210. Since the web server may be hosting multiple hosted sites, it must respond to transactions for each hosted site.

In a web farm configuration, the Service-Level Agreement may depend upon the performance of numerous elements of a server site at the Internet address 116, including the front end server 108, the LAN 110, and each back end server 210A, 210B, 210C,

201D that hosts a virtual site. Therefore, performance and operation of these elements are measured, correlated and reported. Each server 210 keeps a log file of its own activities, including time-stamped records of performance, transactions, and errors. The log file format is generated using modified IIS Server Application Program Interface (ISAPI) filters on Microsoft® Internet Information Server (IIS), or Netscape® Server Application Program Interface (NSAPI) plug-ins on Netscape servers, or extensions for Apache™ servers. A typical log file contains the additional information pertaining to the client IP address, username, date, time, service, server name, server IP address, processing time, bytes sent, bytes received, service status, operation, target URL, User Agent, referrer parameters, SMserver, Smvirtual site, and cookie. Log files created on each server in the web farm also include a time stamp to measure the transfer time, the routing host's name, and the DNS name of the host that serves the request.

In addition, log files can be further extended to include information necessary to report on SLAs that include metrics based upon the additional parameters including the following: the assigned disk space; what the user can access; how the user's request is fulfilled within the system or web farm; user's subscribed level of service or class; transaction; # of requests; download size; file size, file type etc.; time of day, week or month; response time of the back end servers or web farm; and/or how long it takes to complete a specified request or file.

As a further illustration of an embodiment of the invention, Fig. 2 shows a block diagram of a reporting system. Agents 208 on the back-end servers collect and transfer log files from each server 210 to an Accumulator 202. The Accumulator 202 collects log files from the agents 208, combines them, and places the processed results into a database 310. An SLA Reporter System 206 can then query the database 310 to generate SLA reports, according the configuration arranged by an operator's configuration management GUI 503 (Fig. 1). The elements of the illustrative embodiment are now described in further detail.

At least one of the computers on the network of the web farm runs a process designated as the Accumulator 202 (See Fig. 2). The Accumulator 202 collects data from various back end web servers 210 and generates a file in an intermediate format. The semantics in this file support SLA reporting according to the invention. For example, the

Accumulator runs as a NT Service on NT. Intelligent agents 208, implemented in software, are deployed on back-end web servers 210 and configured to work with the Accumulator 202. Based on a user configurable polling interval, the Accumulator 202 “wakes up” and contacts 207 the agents 208 on backend web servers 210. As a result of this, each agent opens its server log files, and transfers 209 the information over to the Accumulator 202.

As shown in Fig. 2, each agent 208 maintains a context 212 of when it was last contacted, which log file was transferred, and how many bytes of the log were transferred. The first time an agent is contacted by the Accumulator 202, the agent will create all the context information and transfer the entire log file. The next time around when the agent is contacted, and the same log file is requested, the agent will determine the difference from the data already transferred and will only transfer the “delta” from the previous transfer. The file transfer 209 from each agent 208 to the Accumulator 202 is made using any well known file transfer protocol, such as FTP over TCP/IP.

After all agents 208 in the backend servers finish their log file transfer 209, Accumulator 202 then combines the log files into a single file 306. The Accumulator 202 then invokes a log parser routine 308 to parse this combined log file 306, recognize the semantic types and upload them into a data repository 310. After the Accumulator 202 parses and updates the data repository 310, it will go back to sleep until the next timed update interval. Updates can also be triggered by selected events, such as trouble signals, or manual intervention.

In the flow chart of Fig. 3, the processing of the Accumulator 202 is provided in greater detail. At step 402, a scheduler triggers the Accumulator to collect the log files 302 from the agents 208. At step 404 the Accumulator obtains a list of the servers 210 having agents 208 to contact. Steps 406 and 408 start an agent contact thread and test to see if the thread is started. If the thread is not started, then the decision at step 408 results in an error condition 410. Once a thread is started, step 412 requests the current log file from each listed agent. As described hereinbefore, each agent 208 maintains its own context file 212 to determine which information is sent to the Accumulator 202. At step 414 each listed agent 208 transfers its corresponding log file 302 to the Accumulator 202. The Accumulator process then waits for the completion of transfer of each log file at

steps 416 and 418. When all listed agents have transferred their files 302, the Accumulator executes step 420 in which the log files are consolidated into a single file 306. Then the database (Db) 310 is updated.

The database 310 provides the information necessary for responding to queries such as from an SLA reporter 502, illustrated in Fig. 4. The SLA Reporter 206 is a process which generates reports. The reporter initially accesses administratively defined SLA parameters or profiles 802. The reporter assembles queries 804 to get actual SLA parameters/measurements from the database. A connection 806 is made to the database. If successful 808, queries are run and the report is generated 812. The report is formatted and delivered 414 in accordance with user-specified options 816. The SLA Reporter will contain queries that will contrast the agreed SLA parameters against the actual measured service level. SLA metrics are provided per virtual site, URL, transaction, content type, file type, source IP, and user class. Within a virtual site, definitions for classes are also created. Each class implies certain agreed-to SLA metrics. For example, a "gold" class may be defined to require higher performance than a "silver" class. A hosted site is defined with no more than one class. Similarly, a User is defined with no more than one class. An operator defines the acceptance parameters of a class by entering the class name, and assigning servers to service that class. Servers may also be re-assigned by other web-farm entities as performance needs indicate. In addition, a class is defined by way of its subscribed error rate, response time, and bandwidth parameters. As an added feature, if a configuration operator attempts to delete a class definition from the database, the system will first check to see if the class is defined for any host, and will prevent deletion of assigned classes.

An SLA Configuration GUI 503 allows the operator to define SLA metrics. In effect, the SLA requirements are programmed into the report generator. The SLA reporter can be configured to provide information according to an Internet IP address (i.e., a "user"), or according to the hosted site name, URL, transaction, content types or file type. This permits reports for either type of SLA. An operator uses a configuration interface to add SLA parameters for any of the foregoing parameters to the database. A User configuration includes: user name, user IP address, subscribed response time, subscribed error rate, and subscribed bandwidth. A Host configuration includes: Host

name and subscribed bandwidth. As an added feature, configuration parameters may also be imported in batch form from a file selected by the operator using a file browsing feature of the GUI. A URL configuration includes file name response time and error data. A transaction configuration includes a group of URLs, subscribed response time and subscribed error rate. A content type configuration includes type of content and subscribed error rate. A file type configuration includes type of files, subscribed error rate, and subscribed bandwidth.

The operator then specifies a "profile" of information to be included in an SLA report. There may be multiple profiles defined and independently named. The GUI permits creating, editing, copying and deleting profiles. Each profile may contain an indication of the database from which SLA reports are to be compiled, and the name of a hosted site to be reported. A home page and web server URL are also defined in the profile, in order to provide meaningful reports. Various filters can be defined and then used to include or exclude information from the SLA reports generated by the profile.

SLA reports are comprised of queries to the database. The queries are based upon specification of the URL, the back-end virtual site, the back-end server, the class, or the user. They may also be based upon the total bandwidth or the bandwidth used by each backend server. The operator defines a report based upon information supplied in response to the desired queries.

Figure 5 illustrates a sample SLA report generated by an embodiment of the present invention. Fig. 5A shows a sample report of service-oriented information accumulated for a series of URLs. The service-oriented information in the sample includes: number of requests for the web page, visits, the subscribed and round-trip delay, the subscribed and actual error rate, and the percent deviation between each set of subscribed and actual results. Similarly, Fig. 5B illustrates a sample SLA report based on the name of a virtual site in a web farm. Similarly, Fig. 5C illustrates a sample SLA report based upon each name of a back-end server in the web farm. The service-oriented information is as described above. Fig. 5D illustrates a sample SLA report for bandwidth of the site, comparing the subscribed to the actual. Fig. 5E shows a sample SLA report for bandwidth utilized by each virtual site named in the web farm.

Fig. 5F illustrates a class-based SLA report, showing the selected performance parameters of several defined classes. Fig. 5G shows a sample SLA report for a list of users, comparing the subscribed versus actual performance, and the percent deviation for each user.

5 It should be appreciated that various other permutations of information can be developed and presented in report form according to the invention.

Further, it should be appreciated that although the illustrative embodiment herein is described in the context of a Windows NT platform, other platforms could provide the basis for implementation of the functional components described herein, such as Solaris,  
10 Linux, other UNIX variations or the like.

Although the functionality described herein is particularly implemented and divided between the Accumulator and the back-end servers/ agents, it should be appreciated that the functionality could be alternatively divided and implemented. For example, certain functionality on the Accumulator could be implemented on the servers  
15 and vice versa. Similarly, although the functionality is generally described implemented as software processes, it should be appreciated that such functionality can be alternatively implemented as hardware, firmware and/or any of various combinations of hardware, software, and firmware.

While particular sets of parameters, i.e. server attributes, are reported to the router  
20 in the implementation described herein, it should be appreciated that parameters other than those described could be reported as a function of the server information available, such as the types of files or size of files, or the like.

Although the invention is shown and described with respect to an illustrative embodiment thereof, it should be appreciated that the foregoing and various other  
25 changes, omissions, and additions in the form and detail thereof could be implemented without changing the underlying invention.




What is claimed is:

1. A method for evaluating real-time compliance with a service-level agreement, said method comprising the steps of:
  - defining a set of parameters to be measured;
  - 5 defining acceptance levels for each parameter in said set of parameters;
  - collecting real-time information related to measurement of said parameters;
  - comparing said acceptance levels to said real-time information; and
  - generating a report.
- 10 2. The method of claim 1 in which said step of defining further includes the steps of
  - providing a format in which a set of servers will provide information to be measured; and
  - implementing a means for collecting said information.
- 15 3. The method of claim 1 in which said step of defining acceptance levels comprises the step of generating a database entry for each service commitment element of a service-level agreement.

## Abstract

A method and apparatus is provided for generating, collecting, and manipulating useful information for validating or defining SLAs of web servers on a network. Web  
5 servers comprising a web farm on the network are adapted for logging detailed runtime information regarding user transactions and performance parameters. An Accumulator device interacts with intelligent agents to collect and combine their log files, process the combined file and post information into a database. An operator enters committed performance parameters into an SLA Reporter system according to classes of users,  
10 classes of web sites being hosted on the web servers, classes of URLs, transaction, content and file type. When compared with the database, processing of SLA reports indicate how well the parameters of the SLAs are being met for users, web sites, classes, URL's and transactions, or other measurable elements. By generating, collecting, combining and processing in this manner, application-specific performance can be  
15 quickly and automatically evaluated with respect to parameters related to user satisfaction and detailed signals can be issued for cases in which remedial steps should be undertaken.

Fig. 1



The diagram illustrates a system architecture for log file processing. It features two agents, Agent1 and Agent2, each with a 'Context Info.' box. The agents interact with a central 'SLA Reporter System' (206) and a database (310). The process involves four steps: 1. GetLogFile, 2. Ftp Transfer Log File, 3. Combine Log Files, and 4. Parse and stuff into a Db. Handwritten labels like 210A, 208A, 207A, 209A, 212A, 210B, 208B, 207B, 209B, 212B, 202, 302A, 302B, 306, 308, and 310 are present.

Fig 2

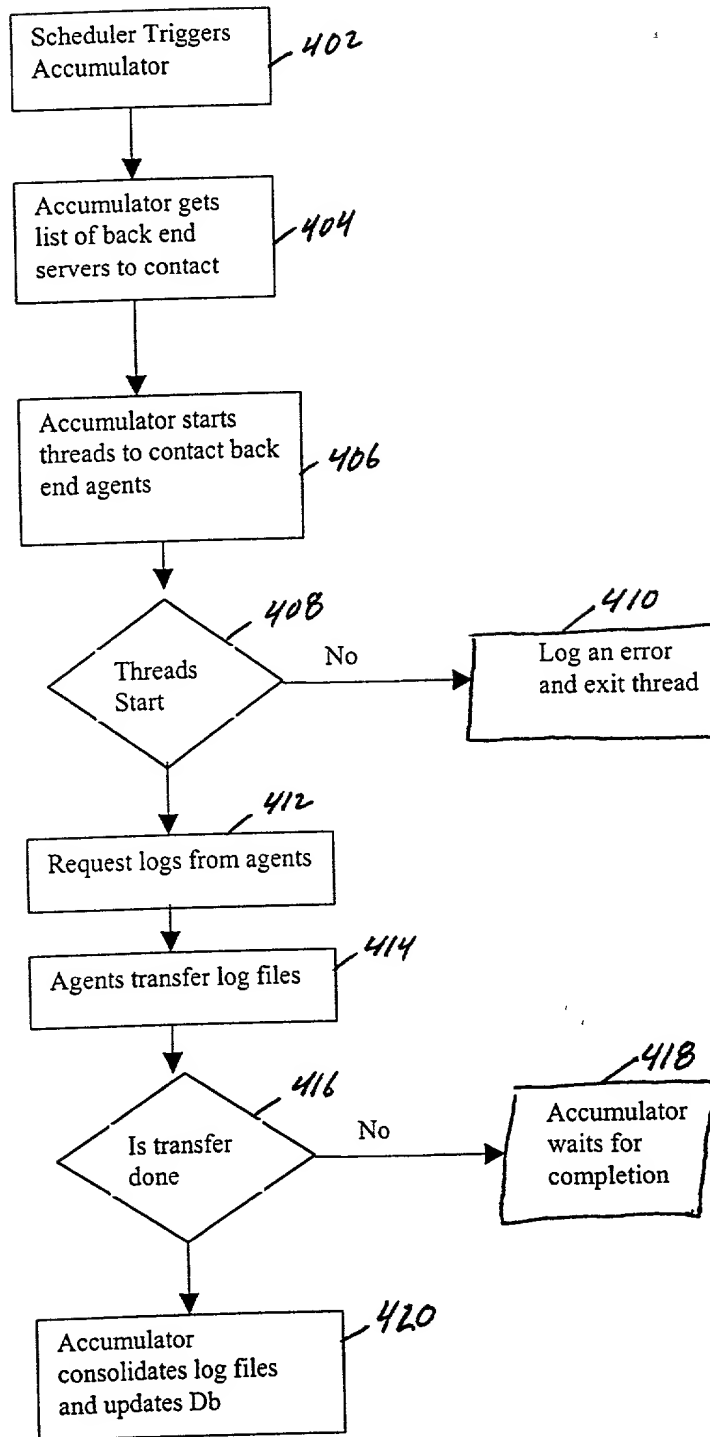
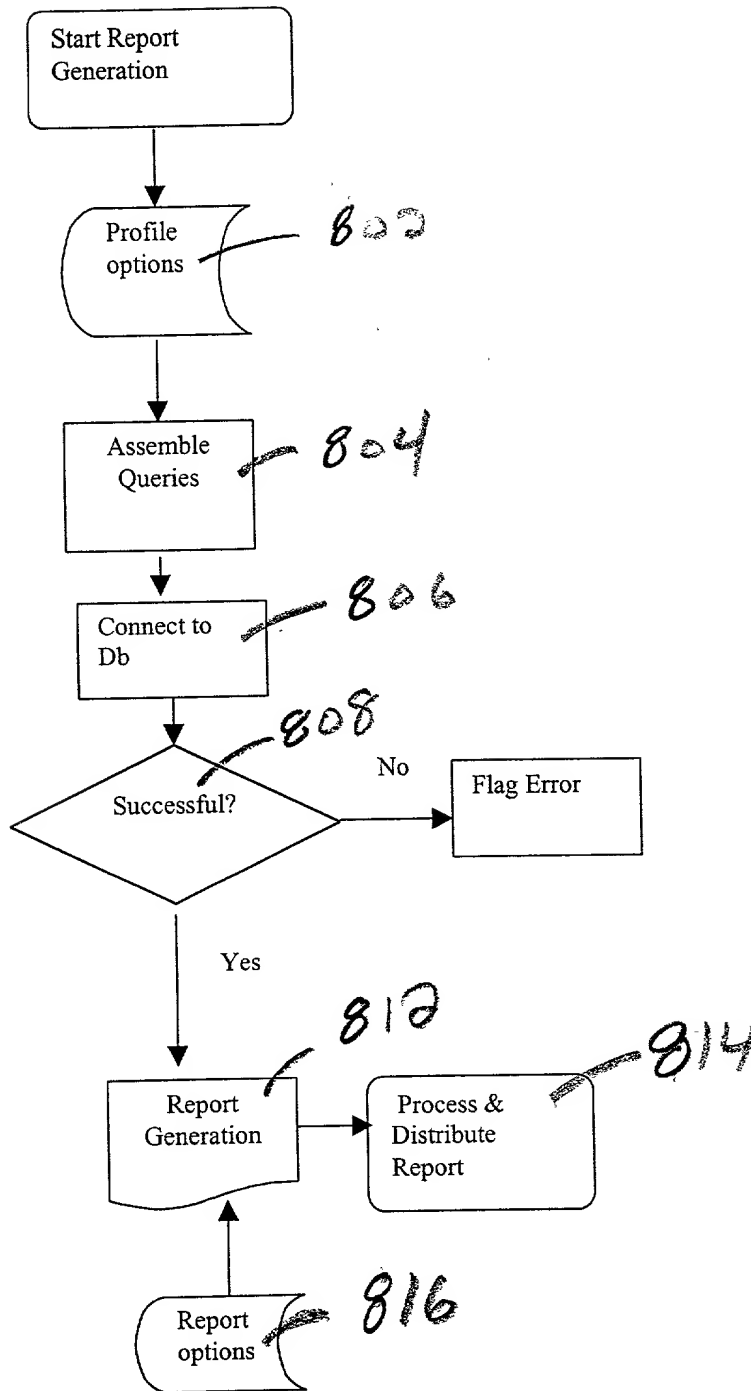


Fig. 3

Fig. 4



URL Based SLA Query URL Based SLA Query

Fig. 5A

URL Based SLA Query									
	Page Name	Requests	Visits	Subscribed Avg.RT (Sec)	Actual Avg. RT (Sec)	% Deviation	Subscribed ErrorRate	Actual Error Rate	% Deviation
1	Webmanage HomePage /homepage.html	62	33	5	5	0	3	1	33.3
2	Default Page /index.html	36	19	5	8	-60	3	3	0
3	Product Details of Webmanage /products/products.html	12	6	5	6	-20	3	2	+33.3
4	About Webmanage /company/aboutus.html	4	2	5	7	-40	3	0	+100
5	Main page /main.html	4	2	5	3	+40	8	0	+100
6	Enroll for Product Evaluation!! /products/evaluation.html	4	2	5	4	+20	3	2	+33.3
	Total	92	64						

Back-end Virtual Site Based SLA Back-end Virtual Site Based SLA

Fig. 5B

Back-end Virtual Site Based SLA							
	Virtual Site Name/IP	Subscribed Avg.RT (Sec)	Actual Avg. RT (Sec)	% Deviation	Subscribed ErrorRate	Actual Error Rate	% Deviation
1	Ws1.bu.edu	5	1	+80	3	2	+33.3
2	Ws1.wmt.edu	5	1	+80	3	2	+33.3
3	Ws2.bu.edu	5	3	+40	3	0	+100
4	Ws2.wmt.com	5	4	+20	3	1	+60

Back-end Server Based SLA Back-end Server Based SLA

Fig. 5C

Back-end Server Based SLA								
	Server Name/IP	Class	Subscribed Avg.RT (Sec)	Actual Avg. RT (Sec)	% Deviation	Subscribed ErrorRate	Actual Error Rate	% Deviation
1	Server1	Platinum	5	1	+80	3	2	+33.3
2	Server2	Platinum	5	8	-60	6	10	-66.6

Total Bandwidth Total Bandwidth

Fig. 5D

Total Bandwidth			
	Bandwith Subscribed(KB)	Actual Bandwidth (KB)	% Deviation
1	9000	7900	+12.2

Bandwidth Utilized by each Backend Server Bandwidth Utilized by each Backend Server **Fig. 5E**

Bandwidth Utilized by each Backend Server		
	Virtual SiteName/IP	Backend Web Server Bandwidth (KB)
1	Ws1.wmt.com	Server1 1200
2	Ws1.bu.edu	Server1 2300
3	Ws2.wmt.com	Server2 3400
4	Ws2.bu.edu	Server2 1000
	Total	7900

Class Based SLA Class Based SLA

**Fig. 5F**

Class Based SLA							
	Class	Subscribed Avg.RT (Sec)	Actual Avg. RT (Sec)	% Deviation	Subscribed ErrorRate	Actual ErrorRate	% Deviation
1	Platinum	5	3	+40	3	0	+100
2	Gold	5	4	+20	3	4	-20
3	Silver	5	5	0	3	2	+33.3

User Based SLA Query User Based SLA Query most popular visitors to your site.

**Fig. 5G**

User Based SLA Query							
	User Name/IP	Class	Subscribed Avg.RT (Sec)	Actual Avg. RT (Sec)	% Deviation	Subscribed Error Rate	Actual Error Rate % Deviation
1	IBM	Gold	5	2	+60	3	1 +66.6
2	Microsoft	Gold	5	6	-20	6	3 +50
3	INTEL	Platinum	9	5	+44.4	8	5 +37.5



Express Mail Label No. **KE 633 278 765 US**

Page 1 of 3

Docket No.  
20496/2

## Declaration and Power of Attorney For Patent Application

### English Language Declaration

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

**METHOD AND APPARATUS FOR PROVISIONING, ASSIGNMENT, MEASUREMENT AND AUDIT SERVICES PROVIDED TO CLUSTERED NETWORK SITE USERS, AND TO GENERATE SERVICE LEVEL AGREEMENT REPORTS.**

the specification of which

(check one)

☒ is attached hereto.

☐ was filed on \_\_\_\_\_ as United States Application No. or PCT International Application Number \_\_\_\_\_ and was amended on \_\_\_\_\_

(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d) or Section 365(b) of any foreign application(s) for patent or inventor's certificate, or Section 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate or PCT International application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s)

Priority Not Claimed

_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<input type="checkbox"/>
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<input type="checkbox"/>
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<input type="checkbox"/>

0045159113050

I hereby claim the benefit under 35 U.S.C. Section 119(e) of any United States provisional application(s) listed below:

<u>60/110,350</u>	<u>DECEMBER 1, 1998</u>
(Application Serial No.)	(Filing Date)
 <u>                    </u>	 <u>                    </u>
(Application Serial No.)	(Filing Date)
 <u>                    </u>	 <u>                    </u>
(Application Serial No.)	(Filing Date)

I hereby claim the benefit under 35 U. S. C. Section 120 of any United States application(s), or Section 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. Section 112, I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, C. F. R., Section 1.56 which became available between the filing date of the prior application and the national or PCT International filing date of this application:

<u>                    </u>	<u>                    </u>	<u>                    </u>
(Application Serial No.)	(Filing Date)	(Status)
		(patented, pending, abandoned)
 <u>                    </u>	 <u>                    </u>	 <u>                    </u>
(Application Serial No.)	(Filing Date)	(Status)
		(patented, pending, abandoned)
 <u>                    </u>	 <u>                    </u>	 <u>                    </u>
(Application Serial No.)	(Filing Date)	(Status)
		(patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

03451446

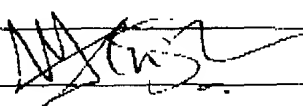
**POWER OF ATTORNEY:** As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. (list name and registration number)

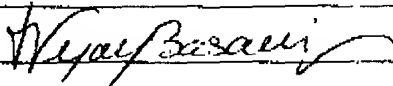
**Brian L. Michaelis (Reg. No. 34,221)**

Send Correspondence to: **Brian L. Michaelis, Esq.**  
**Brown Rudnick Freed & Gesmer - 18th floor**  
**One Financial Center**  
**Boston, MA 02111**

Direct Telephone Calls to: (name and telephone number)

**Brian L. Michaelis (617) 856-8369**

Full name of sole or first inventor	<b>Krishna MANGIPUDI</b>
Sole or first inventor's signature	
Date	<b>11/30/99</b>
Residence	<b>NASHUA, NH</b>
Citizenship	<b>U.S.</b>
Post Office Address	<b>5 Decatur Drive</b>
	<b>Nashua, NH 03062</b>

Full name of second inventor, if any	<b>Vijay BASANI</b>
Second inventor's signature	
Date	<b>11/30/99</b>
Residence	<b>NASHUA, NH</b>
Citizenship	<b>Indian</b>
Post Office Address	<b>26 Kessler Farm Drive, #418</b>
	<b>Nashua, NH 03063</b>

Docket No.  
20496/2

# Declaration and Power of Attorney For Patent Application

## English Language Declaration

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name,

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

**METHOD AND APPARATUS FOR PROVISIONING, ASSIGNMENT, MEASUREMENT AND AUDIT SERVICES PROVIDED TO CLUSTERED NETWORK SITE USERS, AND TO GENERATE SERVICE LEVEL AGREEMENT REPORTS.**

the specification of which

(check one)

☒ is attached hereto.

☐ was filed on \_\_\_\_\_ as United States Application No. or PCT International Application Number \_\_\_\_\_ and was amended on \_\_\_\_\_

(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d) or Section 365(b) of any foreign application(s) for patent or inventor's certificate, or Section 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate or PCT International application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s)

Priority Not Claimed

_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<input type="checkbox"/>
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<input type="checkbox"/>
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<input type="checkbox"/>

<u>60/110,350</u>	<u>DECEMBER 1, 1998</u>
(Application Serial No.)	(Filing Date)
<u> </u>	<u> </u>
(Application Serial No.)	(Filing Date)
<u> </u>	<u> </u>
(Application Serial No.)	(Filing Date)

(Application Serial No.)	(Filing Date)	(Status) (patented, pending, abandoned)
(Application Serial No.)	(Filing Date)	(Status) (patented, pending, abandoned)
(Application Serial No.)	(Filing Date)	(Status) (patented, pending, abandoned)

Patent and Trademark Office-U.S. DEPARTMENT OF COMMERCE

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. *(list name and registration number)*

**Brian L. Michaelis (Reg. No. 34,221)**

Send Correspondence to: **Brian L. Michaelis, Esq.**  
**Brown Rudnick Freed & Gesmer - 18th floor**  
**One Financial Center**  
**Boston, MA 02111**

Direct Telephone Calls to: *(name and telephone number)*  
**Brian L. Michaelis (617) 856-8369**

Full name of sole or first inventor <b>Krishna MANGIPUDI</b>	
Sole or first inventor's signature	Date
Residence <b>NASHUA, NH</b>	
Citizenship <b>U.S.</b>	
Post Office Address <b>5 Decatur Drive</b>	
<b>Nashua, NH 03062</b>	

Full name of second inventor, if any <b>Vijay BASANI</b>	
Second inventor's signature	Date
Residence <b>NASHUA, NH</b>	
Citizenship <b>Indian</b>	
Post Office Address <b>26 Kessler Farm Drive, #418</b>	
<b>Nashua, NH 03063</b>	